

Claims

We claim:

- 1 1. A system for organizing and accessing a database, the system comprising:
2 a primary B+tree index;
3 a secondary B+tree index;
4 a plurality of mapping table row identifiers stored in the secondary B+tree index; and
5 a plurality of database addresses for leaf blocks of the primary B+tree index
6 corresponding to mapping table rows.
- 1 2. The system according to claim 1, wherein 4 bytes of the database addresses are stored
2 in the secondary B+tree index.
- 1 3. The system according to claim 1, wherein the database addresses correspond to a
2 guess-database address.
- 1 4. The system according to claim 1, further comprising:
2 a guess-database address quality statistic for the secondary index and a guess-database
3 address quality statistic for the mapping table, both statistics being operable to assess guess-
4 database address quality.
- 1 5. A method for managing a database system, the method comprising:

creating a secondary index for a B+tree structure, wherein the secondary index structure comprises a plurality of rows each comprising an index key value, a mapping table rowid value and a guess-database address value.

6. The method according to claim 5, further comprising:

inserting a row of the secondary index structure, wherein inserting the row comprises inserting a row comprising an index key value, a mapping table rowid value and a guess-database address value.

7. The method according to claim 5, further comprising:

deleting a row of the secondary index, wherein deleting the row comprises locating a row comprising an index key value and a mapping table row identifier and deleting the row.

8. The method according to claim 5, further comprising:

updating the secondary index, wherein updating the secondary index comprises locating a row of the secondary index comprising an old index key and a mapping table row identifier, deleting the row and inserting in the row a new index key value, a mapping table row identifier and a guess database address.

9. The method according to claim 5, further comprising carrying out a query utilizing the

secondary index, wherein carrying out the query comprises:

utilizing a guess-database address stored as part of a secondary index row to find a row in the primary B+tree structure;

5 obtaining a target database block from the row in the primary B+tree structure;
6 searching the target database block for a row that contains a mapping table row identifier
7 that is the same as a mapping table row identifier stored in the secondary index row; and
8 if the row in the database block matches the target database block, then the correct row in
9 the database has been located and the query is completed.

1 10. The method according to claim 9, wherein if the row in the database block does not
2 match the target database block, carrying out the query further comprises:
3 accessing the mapping table row stored in the secondary index row;
4 utilizing a guess-database address stored in the mapping table row to access a target block
5 of the database;
6 searching the target block for a primary key that matches a primary key stored in the
7 mapping table row; and
8 if the primary key is found, then the query is completed.

1 11. The method according to claim 10, wherein if the primary key is not located carrying
2 out the query further comprises:
3 traversing the primary B+tree structure utilizing the primary key value from the mapping
4 table row to identify the database address to complete the query.

1 12. The method according to claim 11, further comprising:
2 maintaining a guess-database address quality statistic for the secondary index;
3 maintaining a guess-database address quality statistic for the mapping table;

4 utilizing the statistics to assess guess-database address quality; and
5 carrying out the query based upon guess-database quality in the secondary index and
6 mapping table.

1 13. The method according to claim 12, further comprising:
2 estimating guess-database address quality;
3 estimating the cost of the query based upon the estimated guess-database address quality;
4 and
5 carrying out the query starting with an index structure with the lowest estimated cost.

1 14. The method according to claim 5, wherein only 4 bytes of the guess-database address
2 value are stored in the secondary index row.

1 15. A computer program product for performing a process of managing a database system,
2 comprising:
3 a computer readable medium; and
4 computer program instructions, recorded on the computer readable medium, executable
5 by a processor, for performing the steps of:
6 creating a secondary index for a B+tree structure, wherein the secondary index structure
7 comprises a plurality of rows each comprising an index key value, a mapping table rowid value
8 and a guess-database address value.

1 16. A system for performing a database management process, comprising

2 a processor operable to execute computer program instructions; and
3 a memory operable to store computer program instructions executable by the processor,
4 for performing the steps of:
5 creating a secondary index for a B+tree structure, wherein the secondary index structure
6 comprises a plurality of rows each comprising an index key value, a mapping table rowid value
7 and a guess-database address value.